

Online Vision- and Action-Based Object Classification Using Both Symbolic and Subsymbolic Knowledge Representations

Laura Steinert* Jens Hoefinghoff* Josef Pauli*

* Universität Duisburg-Essen, Fakultät für Ingenieurwissenschaften,
Bismarckstrasse 90, 47057 Duisburg (e-mail:
laura.steinert@stud.uni-due.de; jens.hoefinghoff@uni-due.de;
josef.pauli@uni-due.de)

Abstract: If a robot is supposed to roam an environment and interact with objects, it is often necessary to know all possible objects in advance, so that a database with models of all objects can be generated for visual identification. However, this constraint cannot always be fulfilled. Due to that reason, a model based object recognition cannot be used to guide the robots interactions. Therefore, this paper proposes a system that analyzes features of encountered objects and then uses these features to compare unknown objects to already known ones. From the resulting similarity appropriate actions can be derived. Moreover, the system enables the robot to learn object categories by grouping similar objects or by splitting existing categories. To represent the knowledge a hybrid form is used, consisting of both symbolic and subsymbolic representations.

1. INTRODUCTION

At the beginning of the 21st century robots start to be part of the society and act in e.g. private households. However, performing tasks in a strictly controlled laboratory and performing the same tasks in the *real* world are two completely different things. Most of all, the environment is much more unpredictable, dynamic and complex. Yet when robots are supposed to help humans physically in their everyday life, this problem needs to be approached.

In order to interact with their environment robots often need a database for visual identification that contains models of all objects they can encounter. However, for such unspecified environments as human households a database containing models of all possible objects is not feasible. However, if a robot is supposed to learn patterns of interactions with its environment, object identification is essential. To overcome this problem, a system is proposed that is able to learn objects online based on abstract features. Additionally, it is able to group these objects into categories to predict suitable interaction patterns for new objects by comparing them to its learned knowledge base.

The problem of knowledge generalization as well as analogy finding is hardly new and has been studied before. An example system for analogy making is *COPYCAT* by Mitchell [2001]. *COBWEB* by Fisher [1987] on the other hand is an example for knowledge generalization.

Yet it is the ability to generalize and to make analogies that seem to be at the core of human intelligence. As such it is a subject that greatly concerns artificial cognitive systems. The presented approach consists of a hybrid knowledge representation using both symbolic and subsymbolic representations. Additionally, it takes the so called *symbol grounding problem* by Harnad [1990] into account. The

presented system has been evaluated using a NAO robot of the Aldebaran company.

At the beginning of this paper the terms symbolic, subsymbolic and hybrid knowledge representation as well as the symbol grounding problem are defined. Afterwards, an example scenario is presented that is used in the remaining paper. This is followed by a section that describes the actual framework, first the subsymbolic, then the symbolic parts. At the end of the paper, the evaluation of the framework and the results are discussed before a conclusion is given.

2. KNOWLEDGE REPRESENTATION AND SYMBOL GROUNDING

The decision on a specific knowledge representation is always influenced by the task to be solved. When using a *symbolic knowledge representation*, the knowledge is encoded explicitly. Examples for such representations are graphs. Symbolic representations are well suited to represent relations between objects (cf. French [2002]), for instance to model natural language (cf. Harnad [1990]).

Subsymbolic representations on the other hand encode knowledge implicitly, as it is done in neural networks (cf. Luger [2001]). Such a representation can be good for learning motoric tasks or for analyzing sensory data (cf. Harnad [1990]).

Both symbolic and subsymbolic representations seem to be used by the human brain. Therefore, artificial intelligence should incorporate both representations and combine them into a *hybrid representation* (cf. Kokinov [1994]). This view has started to be incorporated into newer artificial intelligence systems (cf. Harnad [1990], Mitchell [2001]). However, since knowledge represented

symbolically can be very abstract, the risk is high that it is no longer connected to the underlying sensory data and is therefore ungrounded. This leads to the symbol grounding problem.

Stevan Harnad [1990] coined the term *symbol grounding problem* by questioning whether a robot really knows what it seems to know. This directly refers to the problem of symbolically represented knowledge being isolated from sensory data. If a robot knew that *dogs can be petted*, it were of hardly any practical use if the robot were not able to recognize a dog standing in front of it.

The presented approach tries to overcome this problem by embedding the sensory information directly into the knowledge base.

3. EXAMPLE SCENARIO

Before the actual framework is presented, an example scenario is given that is used throughout the paper. The robot is presented apples of either green, red or brown color sequentially, as well as wooden rectangular toy blocks that are either green, red or yellow. It is then asked to verbally sort the objects into a toy box, a fruit basket and a rubbish bin. After every action the robot is either positively or negatively rewarded. For every encounter with an object only one action can be chosen. If the robot sorts green or red apples into the fruit basket, it receives a positive reward. The same holds for sorting wooden toy blocks into the toy box and for putting brown apples into the rubbish bin. For all other actions the robot receives a negative reward. At the beginning of the scenario the robot has no knowledge regarding any of the objects.

The expectation is that the robot will create three object categories – one for both green and red apples, one for brown apples and one for all wooden toy blocks.

This scenario seems to be predestined to be solved by a neural network. But while an artificial neural network often needs to be trained anew whenever the scenario changes slightly, the presented system can adapt to such changes by itself. An example of such a change is the introduction of a new object category.

4. FRAMEWORK

In this section the framework is presented. First, a brief overview of the system is given and the basic underlying algorithm is explained. After that, another subsection describes the feature extraction, which is one of two major components of the system. A subsection regarding the knowledge management unit, the second major component of the system, follows.

4.1 Overview

The system presented uses a hybrid knowledge representation as explained above. As such it can be divided into two sections: one that uses a symbolic and one that uses a subsymbolic representation of knowledge. Figure 2 shows an overview of the system.

Whenever an object is presented to the robot, an image of the object is taken. Based on this image, features are



Fig. 1. A red wooden toy block is shown to the robot

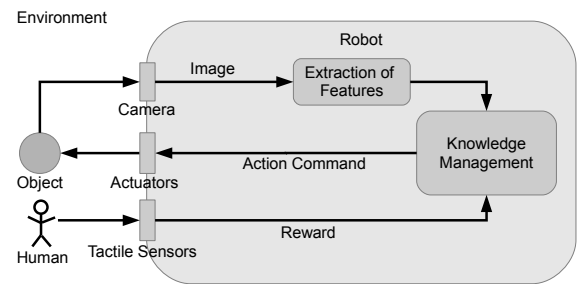


Fig. 2. Overview of the systems parts and their interactions with each other and the environment

extracted – in the case of the example scenario color and basic shape (rectangular or circular). The feature extraction components – one for each feature – can be multithreaded and independent of each other. In these components a subsymbolic representation is used. Furthermore, the usage of other sensors to extract the feature set is also possible.

Figure 1 shows the robot while performing the example scenario. At the displayed time it is shown a red wooden toy block. The foreground of the image shows other objects used in the scenario.

After the feature set is computed, it is relayed to the knowledge management, which uses a symbolic representation for its knowledge base. The robot first checks whether an object that fits the currently perceived feature description has been encountered before. If this is the case, an action based on the previous experience with this object is chosen. If this object has not been seen before, it is compared to all known object categories. Depending on the similarity to other objects, an appropriate action is chosen. In both cases the chosen action is performed. Afterwards the robot is given a reward for the chosen action by a human supervisor. In the implemented version of the system this was accomplished by using the tactile sensors of the robot. The reward is then used to update the knowledge base, e.g. by combining or dividing object categories.

The two following subsections describe in greater detail the symbolic and the subsymbolic parts of the system.



Fig. 3. From left to right: The original image (converted to black and white), the perceived object silhouette and the quartered bounding box

4.2 Feature extraction

The feature extraction components use a subsymbolic knowledge representation. In the case of the example scenario the features *color* and *basic shape* have been extracted from images using neural networks. This feature set alone is of course not sufficient for a robot to interact in a human household. Yet it is sufficient to evaluate the system.

For each of the two features to be extracted, a separate multilayer perceptron with one hidden layer is used. Since the example scenario is executed under natural lighting conditions, color variations stemming from changing light conditions are to be expected. In order to minimize this effect, the L*a*b color space is used and the L-channel – which contains the luminacy data – is omitted.

Yet before the color is analyzed, a preprocess identifies the contour of the object in the image. Afterwards, a bounding box is formed around the object. Figure 3 shows a binary image of an object, the silhouette of the detected object and a quartered bounding box superimposed on the silhouette. For every quarter of the bounding box, the average color values of the pixels inside the object contour in the original colored image are then computed.

As an example, consider the following matrix containing the object's color values of the b-channel of the upper right quarter of the bounding box. The background of the image, i.e. everything of the image that does not depict the object, was replaced by zeros. Therefore, all non-zero values encode pixels of the object.

```

0  0  0  0  0
130 0  0  0  0
134 137 0  0  0
138 135 139 0  0
140 138 140 0  0

```

For each of such quarters and for both the a- and b-channel values the average of all pixels with values greater than zero is computed. In the example case this is approximately 137.

The resulting eight values are used as an input for the multilayer perceptron, classifying the color to be red, green, yellow or brown. The classification values are then converted to percentages.

The used multilayer perceptron has one hidden layer consisting of ten nodes and was trained with the *RPROP* (resilient backpropagation) algorithm.

In order to analyze the form of an object, the object contour is extracted first. Afterwards, another bounding box




Fig. 4. The original image (left) and the perceived object silhouette together with the bounding box and minimal enclosing circle

with varying orientation as well as the minimal enclosing circle are computed. Then the area of the bounding box, the circle and the object itself are determined. The relation of the area of the object to that of the bounding box as well as the relation of the area of the object to that of the circle are then used as an input for the multilayer perceptron analyzing the shape. This artificial neural network has one hidden layer with two nodes and was trained with the *RPROP* algorithm, too.

Figure 4 displays the necessary steps before the actual form extraction takes place. The left-hand side of the figure shows the original image. On the right, the object silhouette as well as the bounding box and minimal enclosing circle are shown.

The resultant of the multilayer perceptron is again a vector, containing the classification values for the shapes *rectangular* and *circular*. It is converted to percentages as well.

Both classifiers have been thoroughly tested. The following numbers give the evaluation results for the color test. Hereby, $Y(o)$ yields the classification result with an object o as input and $X(o)$ returns the real feature of the object.

$$\begin{aligned}
P(Y(o) = \text{red} \mid X(o) = \text{red}) &\approx 0.82 \\
P(Y(o) = \text{green} \mid X(o) = \text{green}) &\approx 0.86 \\
P(Y(o) = \text{yellow} \mid X(o) = \text{yellow}) &\approx 0.96 \\
P(Y(o) = \text{brown} \mid X(o) = \text{brown}) &= 1
\end{aligned}$$

$$\begin{aligned}
P(X(o) = \text{red} \mid Y(o) = \text{red}) &= 1 \\
P(X(o) = \text{green} \mid Y(o) = \text{green}) &\approx 0.93 \\
P(X(o) = \text{yellow} \mid Y(o) = \text{yellow}) &= 1 \\
P(X(o) = \text{brown} \mid Y(o) = \text{brown}) &\approx 0.77
\end{aligned}$$

Thus, $P(Y(o) = a \mid X(o) = a)$ gives the conditional probability that the system classifies the color of an object o to be a , under the condition that the color was a . $P(X(o) = a \mid Y(o) = a)$ on the other hand gives the conditional probability that an object has the color a , given that the system classified it as having color a .

As the probabilities show, most misclassifications happened with red objects, which were mostly mistaken for brown ones. The same holds for green objects. This is reflected in the last probability, showing that if the system reported an object to be brown, this was only correct in 77 percent. All together this still gives a color classification that is sufficiently reliable for our purposes.

In the following the conditional probabilities for the shape evaluation are given – the abbreviation *rect.* stands for *rectangular*. As can be noticed, that the reliability is even higher than that of the color classification. Hence, this classification is sufficient for our purpose, too. Both clas-

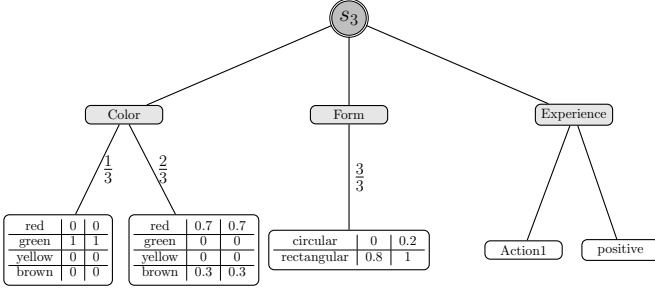


Fig. 5. An example knowledge graph containing data of one object category

sifiers are learned offline, i.e. before the robot is presented objects.

$$\begin{aligned} P(Y(o) = \text{circular} \mid X(o) = \text{circular}) &= 1 \\ P(Y(o) = \text{rect.} \mid X(o) = \text{rect.}) &= 0.97 \end{aligned}$$

$$\begin{aligned} P(X(o) = \text{circular} \mid Y(o) = \text{circular}) &\approx 0.97 \\ P(X(o) = \text{rect.} \mid Y(o) = \text{rect.}) &= 1 \end{aligned}$$

4.3 Knowledge Management

The knowledge management unit is one of two major components of the system. All learned knowledge is administered here, including the analogy making and knowledge generalization. Thus, the actual action selection in response to a presented object is performed in this component, too (cf. figure 2). The first subsection briefly describes how the knowledge is represented. This is followed by an overview of the algorithm that selects actions and evolves the knowledge base.

Knowledge Representation

As has been mentioned before, the knowledge is represented symbolically by a graph. This graph contains the data of all known object categories, including the features linked to each object category and the experiences the agent has made with it as well as its similarity to other object categories.

An example graph containing only one object category can be seen in figure 5. The object category in this example is labeled as s_3 and has the color features

$$[re : 0.7 - 0.7; gr : 0 - 0; ye : 0 - 0; br : 0.3 - 0.3]$$

as well as

$$[re : 0 - 0; gr : 1 - 1; ye : 0 - 0; br : 0 - 0]$$

attached to it. Hereby *re* denotes *red*, *gr* *green*, *ye* *yellow* and *br* denotes *brown*. The entry for each color characteristic – e.g. *red* – is an interval.

The feature intervals connected to an object category are needed to determine whether a perceived object fits into an object category. An object is defined as fitting into a category if for each of its feature vectors an appropriate feature vector can be found whose characteristic intervals

all contain the characteristic values of the object. Please note that this must hold for all features.

In the example graph an object with the color

$$[0.7, 0, 0, 0.3]$$

and the form

$$[0.1, 0.9]$$

would fit into the category s_3 . But an object with the color

$$[0.2, 0, 0, 0.8]$$

and the same form would not fit into the category, since for the color vector no vector associated with the category can be found, such that all characteristics lie within the characteristic intervals of the vector.

The links connecting the feature vectors to the name of the corresponding feature have an occurrence probability.

In the example case this means that so far three objects of the category s_3 have been encountered, of which two were to 70 percent red, the third green in color. Furthermore, objects of this category have a form that has been classified as 0-20 percent circular vs. 80-100 percent rectangular.

As experience a *positive* reward for *Action1* has been obtained. In the example scenario this could be an object category for red and green toy blocks where *Action1* denotes the sorting into the toy box.

Knowledge Management Algorithm

Whenever the system detects an object, the extracted feature vectors are relayed to the knowledge management component. Here it is checked whether the reported object already fits into an object category as described above.

In case the detected object fits into a category, the respective feature vector probabilities as well as the similarity of this object category to all others are updated. Afterwards, an appropriate action is selected. How such a similarity is calculated will be explained shortly.

If the detected object does not fit into any known object category, a new object category is created. Then the similarity of the new category to all other categories is calculated and an appropriate action is chosen. Next the action selection process is described.

Action selection An action is selected based on previous experiences according to the following scheme.

- Choose an action that has a positive experience with the present object category.
- If such an action does not exist, the object categories with a positive similarity to the current category are checked for positive experiences, starting with the most similar category. Nonetheless, the action selected should not have received a negative reward

with either the current or the most similar object category.

- If such an action still does not exist, a random action is chosen from the set of actions that have no negative experience with the current object category.

Instead of using the above presented scheme to determine an action, a random action is sometimes chosen to avoid staying in local maxima – the probability is given by the parameter ρ_{ra} .

ρ_{ra} : Probability of choosing
a random action

Suppose that in the example scenario the robot has only encountered red apples and toy blocks of various colors so far. Now it is presented a green apple. This green apple does not fit into any of the known categories, consequently no experiences exist with it. But by comparing its features to the known categories *red apples* and *toy blocks*, it could detect a high similarity with the *red apples* category. The robot has made positive experiences with sorting red apples into the fruit basket. Therefore, this action is chosen for the green apple, too.

Similarities Between Object Categories In the above description, similarities between object categories have been mentioned multiple times. According to Tversky [1977] not only similar, but also differing features should be taken into account for such a similarity measurement. Luger [2001] further mentions that not all features should be weighted equally. As a consequence, it is proposed that the similarity between two object categories is calculated by summing up the weighted similarities regarding single features as well as the weighted similarity regarding experience.

Therefore, let the following symbols be given:

M	Number of features
s_j, s_k	Two object categories
ω_i	Weight of feature i
ω_e	Weight for the experience
σ_{f_i, s_j, s_k}	Similarity of the object categories s_j and s_k regarding feature f_i
σ_{e, s_j, s_k}	Similarity of the object categories s_j and s_k regarding the experience

Then the *similarity between two object categories* s_j and s_k is computed with equation (1).

$$\sigma_{s_j, s_k} = \sum_{i=1}^M \omega_i \cdot \sigma_{f_i, s_j, s_k} + \omega_e \cdot \sigma_{e, s_j, s_k} \quad (1)$$

The similarity of two object categories regarding a single feature (*feature similarity*) is calculated using equation (2).

$$\sigma_{f, s_j, s_k} = \sum_{c_j \in C_j} \max_{c_k \in C_k} \left((1 - \Delta_{jk}) \cdot P(c_j) \cdot P(c_k) \right) \quad (2)$$

In this equation, C_j and C_k are the sets of associated feature vector intervals of feature f for the object categories s_j and s_k . The cardinality of C_j is always smaller or equal to the cardinality of C_k . If this is not the case, s_j and s_k are interchanged. This guarantees that the similarity of two object categories is not dependent on the order in which they are compared to each other – i.e. s_j compared to s_k produces the same result as s_k compared to s_j . $P(c_i)$ represents the occurrence probability of the feature vector interval c_i for the object category s_i . Δ_{jk} on the other hand encodes the overall distance of the compared feature vector intervals by summing up the shortest difference between the intervals for each characteristic of the feature.

As an example the feature similarity between the following feature interval vectors is calculated.

$$\begin{aligned} [re : 0.7 - 0.7; \text{ gr} : 0 - 0; \text{ ye} : 0.3 - 0.3; \text{ br} : 0.1 - 0.1] \\ [re : 0.6 - 0.8; \text{ gr} : 0 - 0; \text{ ye} : 0 - 0; \text{ br} : 0.3 - 0.4] \end{aligned}$$

The two *red* (re) characteristic intervals have a distance of 0, since they overlap. The same holds for the *green* (gr) characteristic intervals. The intervals for the characteristic *yellow* (ye) on the other hand do not overlap, but have a distance of 0.3; the distance for the *brown* (br) characteristic is $0.3 - 0.1 = 0.2$. By summing up these values up, a value of 0.5 is received for Δ .

The value of Δ has a range of $[0;2]$, whereby smaller values represent a greater similarity. To make this value more intuitive, it is mapped to $[-1;1]$ in which negative values represent a dissimilarity and positive values a similarity. This is achieved by subtracting Δ from 1 (see equation (2)). In order to ensure that the similarity measurement values common feature vector intervals more highly than uncommon ones, the occurrence probabilities of the feature vector intervals are taken into account as well (see equation (2)). These occurrence probabilities are encoded in the knowledge graph, as can be seen in figure 5.

Dynamic Weight Adaption As mentioned before, different features as well as the experience are weighted differently for the similarity calculation. The *weights* – named ω_i and ω_e in equation (1) – all have an initial value of 1 and are dynamically adapted at run time. At all times all weights are greater or equal to zero and the sum of all weights is constant. An adaption of weights occurs in three different cases:

- Two object categories are merged.
- An object category is split into two categories.
- An experience is made for the first time for a given category-action pair and it contradicts the experience of the most similar object category.

In the first case, the two object categories are thought of as very similar. Therefore, the weights should make the detection of similarities of the same type easier. This is achieved by decreasing the weight and therefore the importance of that attribute (feature or experience) in which the two categories differ the most. All other weights are increased. In the second case, the resulting two object categories are apparently not similar although the similar-

ity calculation implies that. Consequently, the weight of the attribute that is the most responsible for their alleged similarity is decreased and all other weights are increased. The same holds for the third case. By how much the weights are decreased or increased at a time is controlled by the parameter δ_{aw} .

δ_{aw} : Controls by how much weights
are in- or decreased

Reward Processing and Knowledge Generalization So far the overall algorithm has chosen an action as a response to a presented object. This action is executed in the next step and a reward is given by a human operator. This experience is then automatically added to the knowledge base. If such an experience already exists in the graph and both – the saved and the newly received reward – are the same, no update is needed. If one of the rewards is *neutral*, the other reward is saved. However, if one reward is *positive* and the other *negative*, this is perceived as a strong indicator that the object does not fit the object category. The currently presented object is then *split* from the object category to form a new category. In any case the similarity values between object categories have to be updated again.

Afterwards as long as object categories exist whose similarity is greater than or equal to a threshold θ_{mc} , these object categories are *merged*. During a merging operation feature vector intervals can also be merged if they have a distance Δ (cf. above) smaller than or equal to another threshold θ_{mf} . After a merging operation, similarity values between object categories have to be recalculated again.

θ_{mc} : Threshold, controls
merging of categories

θ_{mf} : Threshold, controls
merging of feature vector intervals

In the example, the agent has chosen to sort the green apple into the fruit basket. After the execution, it receives a positive reward for this action from a human supervisor. This experience is then entered into the knowledge base. Since no previous experience existed, splitting is not considered. Furthermore, the experience of the most similar object category – *red apples* – with this action is not contradictory, therefore no weight adaption is needed either. Yet the similarities to all other object categories have to be updated to take the new experience into account. As a result, the similarity with the object category *red apples* is larger than the threshold θ_{mc} . Hence, the two object categories are merged. The two merged categories differed the most in the color feature. As a result, the weight of the color feature is decreased by the value of δ_{aw} for the similarity calculation and the weights for color and experience are each increased by the value of $\delta_{aw}/2$. Consequently, the similarities between all object categories have to be updated again. If no other object categories need to be merged, the robot is ready to interact with another object.

Overall, the results of the algorithm are controlled by four parameters. The parameter ρ_{ra} controls how often random actions are chosen. δ_{wa} defines the amount by which the weights for the different attributes are decreased or increased for the similarity calculation. The threshold parameters θ_{mc} and θ_{mf} regulate the merging of categories and feature vectors respectively. How the parameter settings affect the results is discussed in the next section.

5. EVALUATION OF THE PROPOSED FRAMEWORK

The presented system was evaluated with two different scenarios. The first was the example scenario mentioned in section 3. This scenario was also used to analyze the effects the parameters have on the system's performance. The second tested scenario was the so called *Wisconsin Card Sorting Test* (WCST), which will be explained later.

5.1 Example Scenario

To evaluate the system's performance for the example scenario, it was analyzed after how many presented objects the desired object categories had been formed and were not changed anymore. For the example scenario, the desired outcome consists of three object categories: one for both green and red apples, one for toy blocks and one for brown apples. The system was tested using real world objects. However, to prevent feature extraction problems to alter the test, the scenario was additionally tested with simulated feature vectors in two variants: In one variant, the simulated features were 100 percent correct, in another the features were classified correctly but contained uncertainties.

An example of a simulated, 100 percent correct color feature vector is given as follows:

$[re : 1; gr : 0; ye : 0; br : 0]$

An example for a simulated color feature vector containing uncertainties on the other hand is given as:

$[re : 0.8; gr : 0; ye : 0.05; br : 0.15]$

In all three variants of the scenario parameter settings could be found for which the system reached the desired outcome very quickly. A lower bound for the number of steps it takes the system to acquire all requested object categories is of course the number of objects shown to the system before it has seen an object of every category. With some parameter settings the system was able to reach this lower boundary, meaning that it had formed the desired object categories as soon as it had seen at least one object of every category.

The order in which objects were presented to the system was also permuted, yet it influenced the system's performance only slightly. During the real world tests a few misclassifications occurred. However, the system was only temporarily 'confused' and managed to reach the desired outcome none the less. As has been mentioned briefly, the chosen parameters have an influence on the system's




Fig. 6. System performance for different parameter combinations

performance. In order to evaluate the impact of two of the four parameters, the example scenario was tested in all three variants with different parameter settings.

Figure 6 shows the results for variations of the threshold parameter (θ_{mc}) and the weight adaption parameter (δ_{aw}). The other two parameters remained constant during the test with values of $\rho_{ra} = 0$ and $\theta_{mf} = 0.3$. Red crosses indicate that the desired object categories were not reached within the allotted time frame; dots indicate that the desired object categories were formed, the color encoding the number of objects that had to be presented before the state was reached. The encircled area contains parameter combinations which produced good results in all test. It has to be noted that only after approximately 25 objects had been presented to the robot, objects of all three categories had been shown. Therefore, the desired object categories could only be found after 25 steps. As can be seen in figure 6, the best parameter combinations reach this state slightly after 25 steps and are therefore very fast.

Overall the parameter θ_{mc} seems to have a larger influence on the system performance than the δ_{aw} parameter. However, the parameter combinations that produce good results form a large area, meaning that slight parameter variations can be tolerated.

5.2 Wisconsin Card Sorting Test

The system’s relearning abilities were evaluated using the *Wisconsin Card Sorting Test* by Berg [1948]. In this test the subject is shown four cards. One card shows one red triangle, another two green stars. A third card depicts three yellow crosses and a fourth four blue circles. The subject is then given a set of 60 cards, each displaying up to four objects. One a single card all objects have the same shape and the same color from the set of four colors and four shapes mentioned above. The subject is then asked to assign each card to one of the four other cards. This setting is depicted in figure 7.

After every assignment the proband is told if the action was correct. Unknown to the subject, he has to sort the cards either according to color, form or number of shown




Fig. 7. The setting of the *Wisconsin Card Sorting Test*

objects. However, whenever he assigns five cards correctly in a row, the rule which determines whether an assignment is correct is changed. The test ends when the subject assigned nine times five cards in a row correctly. If this has not been accomplished after 60 cards, the cards are mixed and given back to the proband.

This test was performed with human subjects by Berg [1948]. Fifty percent of the subjects managed this test quite well. However, the rest of the test group was not able to complete the test. These subjects especially had problems with relearning. Despite receiving negative rewards after the assignment rule changed, they kept choosing previously correct actions. For these subjects, the test had to be aborted after some time.

The relearning abilities of the proposed system were evaluated by performing this test using simulated feature vectors. Unlike the human subjects, the robot was able to complete the test at all times, yet the time required to relearn was far longer than for the human subjects that completed the test. Additionally, these human subjects were sometimes able to predict the rule changes, thereby allowing them to react more quickly to them. But since the system was not designed to keep track of such changes, the robot was unable to do so. Though when looking at the similarity weights for the different features, it can be seen that at a time directly preceding a rule change, the weight of the feature to use as a sorting rule is quite higher than the other weights. This implies that the system correctly identified the underlying sorting rules.

6. CONCLUSION

Modern artificial intelligence systems should use a hybrid knowledge representation. Furthermore, they should take the symbol grounding problem into account. The presented system fulfills both requirements. It offers a flexible way to learn object categories and behavior patterns in unpredictable environments and allows the robot to form analogies to react to new objects appropriately. The symbol grounding problem is avoided by embedding the sensory information directly into the symbolic knowledge representation.

The system has been tested both in real and in simulated environments. The algorithm’s parameters have been shown to influence the results, yet a large group of param-

ter settings exists for which good results could be achieved in all tested scenarios.

Yet the system also makes assumptions that might be too restrictive. For once it assumes that the similarity between objects is calculated the same way no matter what objects are involved. As a result it is assumed that all objects can be grouped in a similar way. Nonetheless, in real world scenarios this might not be the case. The experience factor for the calculation of the similarity between object categories can balance this, but only at the cost of feature importance. Furthermore, the way object categories are formed heavily depends on the extracted features. If the key for object classification lies in a feature that is not detected, no reasonable object categories can be formed. As such, the applicability of the system needs a feature detection that is as diverse as possible.

While performing the tests, different misclassifications occurred. These errors often led to the formation of isolated object categories that were never merged with other categories. As a result, the graph grew unnecessarily large and the system could no longer reach the desired state of having exactly the predicted object categories. Nonetheless the actions selected by the system were unaffected by these additional object categories. To prevent the long time persistence of erroneous object categories in the future, a decay of knowledge that relies on a single occurrence could be introduced.

Another flaw of the presented system is the reward model which only anticipates *positive*, *negative* and *neutral* rewards. Additionally it assumes that the reward for an object is always the same. To overcome this flaw, an advanced model of rewards such as that of so called *somatic markers* by Hoefinghoff [2012], Hoefinghoff [2013] could be used which would also allow varying rewards.

Besides solving the above mentioned problems and too restrictive assumptions, future work might concentrate on the incorporation of multiple knowledge graphs, one for each task to be performed by the robot. This is necessary since the reward for a single action highly depends on the task that shall be performed. The same holds for the formation of object categories. Imagine different scenarios in which a robot shall choose objects for decoration purposes. The action *choose red apple* might receive a positive reward when decorating a red environment, while the same action for the same object might be punished if a blue environment shall be decorated. To solve this problem, a separate knowledge graph could be formed for every task. But since knowledge from one task might be reused in another task, the different graphs should be interconnected. This means that whenever a new object occurs for a certain task, experiences with this object from a different task could be used.

REFERENCES

- Esta A. Berg. A simple Objective Technique for Measuring Flexibility in Thinking. *Journal of General Psychology*, 39(1): 15–22, 1948.
- Douglas H. Fisher. Knowledge Acquisition Via Incremental Conceptual Clustering. *Machine Learning*, vol. 2 (2): 139–172, 1987.
- Robert M. French. The Computational Modeling of Analogy-making. *Trends in Cognitive Sciences*, G (5): 200–205, 2002.
- Stevan Harnad. The Symbol Grounding Problem. *Physica D*, 42: 335–346, 1990.
- Jens Hoefinghoff and Josef Pauli. Decision Making Based on Somatic Markers. In *Proceedings of the Twenty-Fifth International Florida Artificial Intelligence Research Society Conference*, 163–168. AAAI Press, 2012.
- Jens Hoefinghoff, Laura Steinert and Josef Pauli. An easily adaptable Decision Making Framework based on Somatic Markers on the Nao-Robot. *Kognitive Systeme*, 2013.
- Boicho N. Kokinov. A Hybrid Model of Reasoning by Analogy. In K. J. Holyoak et al., editor, *Advances in Connectionist and Neural Computation Theory*, volume 2, 247 – 318. Ablex, Norwood, 1994.
- George F. Luger. Künstliche Intelligenz. Strategien zur Lösung komplexer Probleme. Pearson Studium, Munich, 2001.
- Melanie Mitchell. Analogy-Making as a Complex Adaptive System. In L. Segel and I. Cohen, editors, *Design Principles for the Immune System and Other Distributed Autonomous Systems*, 335 – 360. Oxford University Press, New York, 2001.
- Amos Tversky. Features of Similarity. *Psychological Review*, 84(4): 327–352, 1977.